ON THE COHOMOLOGY RINGS OF TREE BRAID GROUPS

DANIEL FARLEY AND LUCAS SABALKA

ABSTRACT. Let Γ be a finite connected graph. The (unlabelled) configuration space $UC^n\Gamma$ of n points on Γ is the space of n-element subsets of Γ. The n-strand braid group of Γ, denoted $B_n\Gamma$, is the fundamental group of $UC^n\Gamma$.

We use the methods and results of [10] to get a partial description of the cohomology rings $H^*(B_nT)$, where T is a tree. Our results are then used to prove that B_nT is a right-angled Artin group if and only if T is linear or $n < 4$. This gives a large number of counterexamples to Ghrist's conjecture that braid groups of planar graphs are right-angled Artin groups.

1. INTRODUCTION

If Γ is a finite connected graph and n is a natural number, then the *unlabelled configuration space* of n points on Γ, denoted $UC^n\Gamma$, is the space of n-element subsets of Γ, endowed with the Hausdorff topology. The *labelled configuration space* $C^n\Gamma$ is the space of n-tuples of distinct elements in Γ. The *n-strand braid group of Γ*, denoted $B_n\Gamma$, is the fundamental group of $UC^n\Gamma$; the *n-strand pure braid group of Γ*, $PB_n\Gamma$, is the fundamental group of $C^n\Gamma$.

Various properties of graph braid groups have been established by other authors. Ghrist showed in [14] that the spaces $C^n\Gamma$ are $K(PB_n\Gamma, 1)$s, and that a $K(PB_n\Gamma, 1)$ is homotopy equivalent to a complex of dimension at most k, where k is the number of vertices in Γ of degree at least 3. He also made the following conjecture:

Conjecture 1.1. [1, 14] *The (pure) braid group of any planar graph is a right-angled Artin group.*

Abrams [1] (for all n) and Hu [17] (for the case $n = 2$) introduced a discretized configuration space $D^n\Gamma$, and showed that $C^n\Gamma$ and $D^n\Gamma$ are homotopy equivalent under appropriate hypotheses (which are easy to satisfy). Abrams went on to prove that the universal cover of the space $D^n\Gamma$ is a CAT(0) cubical complex. This implies, in particular, that graph braid groups have solvable word and conjugacy problems [2]. Abrams also showed that $PB_2(K_5)$ and $PB_2(K_{3,3})$ are the fundamental groups of closed surfaces, and thus aren't right angled Artin groups. This is the reason for the word "planar" in Conjecture 1.1. Crisp and Wiest [6] have shown that all graph braid groups embed in right-angled Artin groups, which implies that graph braid groups are linear, bi-orderable, residually finite, and residually nilpotent. Connolly and Doig [5] showed that the braid group of any linear tree is a right-angled Artin group. (A tree T is *linear* if there is an embedded arc which passes through every vertex in T of degree at least 3.)

This paper continues a project begun in [10]. In [10], we used a discrete version of Morse theory (due to Forman [13]) to simplify the configuration spaces $UC^n\Gamma$ within

their homotopy types. Our immediate goal was to settle Conjecture 1.1. We were able to compute presentations $\mathcal{P}(B_nT)$ for all braid groups B_nT, where T is a tree; that is, for all *tree braid groups* ([10], Theorem 5.3). The generators of $\mathcal{P}(B_nT)$ are in one-to-one correspondence with critical 1-cells of $U\mathcal{C}^n\Gamma$ and relators correspond to critical 2-cells. Here "critical" is used in the sense of Forman's discrete Morse theory. In [11] it was shown that $H_i(U\mathcal{C}^nT)$ (equivalently, $H_i(B_nT)$, since $U\mathcal{C}^n\Gamma$ is aspherical for any graph Γ [1], [14]) is a free abelian group of rank equal to the number of critical i-cells in $U\mathcal{C}^nT$. It follows from this that $\mathcal{P}(B_nT)$ has the minimum possible number of generators and relators. We were unable to produce counterexamples to Conjecture 1.1, although the form of the relators in $\mathcal{P}(B_nT)$ made a negative answer seem likely for most trees and most natural numbers n.

Here we get nearly complete information about the mod 2 cohomology rings of tree braid groups. Our results allow us to prove that most tree braid groups are not right-angled Artin groups (see Theorem 5.11). Thus we produce a large number of counterexamples to the version of Conjecture 1.1 in which the word "pure" is omitted. (It is worth noting here that Abrams and Ghrist made the conjecture only for pure braid groups. In this paper, we refer to either version of Conjecture 1.1 as "Ghrist's conjecture". We believe that the analogue of Theorem 5.11 will be true for pure braid groups.)

The argument is as follows. We first compute the cohomology ring of B_4T_{min}, where T_{min} is the minimal nonlinear tree. Our calculation shows that B_4T_{min} is not a right-angled Artin group, since $H^*(B_4T_{min}; \mathbb{Z}/2\mathbb{Z})$ is not the exterior face ring of a flag complex (see Section 5). If T is any nonlinear tree and $n \geq 4$, we embed $U\mathcal{C}^4T_{min}$ into $U\mathcal{C}^nT$. By analyzing the kernel of the map on cohomology, we can conclude that B_nT is also not a right-angled Artin group, since its cohomology ring also fails to be the exterior face ring of a flag complex.

Finally, we note that our description of the mod 2 cohomology rings of tree braid groups is likely to have other applications. For instance, Michael Farber [9], [8] has defined an invariant $TC(X)$ of a topological space X, called the *topological complexity* of X, which is an integer measuring the complexity of motion-planning problems of systems having X as their configuration space. Farber establishes cohomological lower bounds for $TC(X)$ in [9].

This paper is organized as follows. In Section 2 we give a brief description of discrete Morse theory and its applications to computing homology. In Section 3 we describe Morse matchings on the spaces $U\mathcal{C}^n\Gamma$. In Section 4, we give a partial description of the mod 2 cohomology ring of any tree braid group. In Section 5, we use the results of Section 4 and a cohomological argument to determine which tree braid groups are right-angled Artin groups.

We would like to thank Ilya Kapovich and Robert Ghrist for participating in discussions related to this work. We thank Aaron Abrams and Carl Mautner for telling us of some of Mautner's counterexamples to Conjecture 1.1 (in both the pure and regular cases) [18], and for suggesting that cohomological methods might be used to produce counterexamples.

2. Background on Discrete Morse Theory

2.1. **Basic Definitions.** In this subsection, we collect some basic definitions from [10] (see also [3] and [13], which were the original sources for these ideas).

Let X be a finite regular CW complex. Let K denote the set of open cells of X. Let K_p be the set of open p-cells of X. For open cells σ and τ in X, we write $\sigma < \tau$ if $\sigma \neq \tau$ and $\sigma \subseteq \overline{\tau}$, where $\overline{\tau}$ is the closure of τ, and $\sigma \leq \tau$ if $\sigma < \tau$ or $\sigma = \tau$.

A *partial function* from a set A to a set B is a function defined on a subset of A, and having B as its target. A *discrete vector field* W on X is a sequence of partial functions $W_i : K_i \to K_{i+1}$ such that:

(1) Each W_i is injective;
(2) if $W_i(\sigma) = \tau$, then $\sigma < \tau$;
(3) $\mathrm{im}\,(W_i) \cap \mathrm{dom}\,(W_{i+1}) = \emptyset$, where im denotes image and dom denotes domain.

Let W be a discrete vector field on X. A *W-path of dimension p* is a sequence of p-cells $\sigma_0, \sigma_1, \ldots, \sigma_r$ such that if $W(\sigma_i)$ is undefined, then $\sigma_{i+1} = \sigma_i$; otherwise $\sigma_{i+1} \neq \sigma_i$ and $\sigma_{i+1} < W(\sigma_i)$. The W-path is *closed* if $\sigma_r = \sigma_0$, and *non-stationary* if $\sigma_1 \neq \sigma_0$. A discrete vector field W is a *Morse matching* if W has no non-stationary closed paths.

If W is a Morse matching, then a cell $\sigma \in K$ is *redundant* if it is in the domain of W, *collapsible* if it is in the image of W, and *critical* otherwise. Note that any two of these categories are mutually exclusive by condition (3) in the definition of discrete vector field.

The ideas "discrete Morse function" and "Morse matching" are largely equivalent, in a sense that is made precise in [13], pg. 131. In practice, we will always use Morse matchings instead of discrete Morse functions in this paper (as we also did in [10]). A Morse matching is sometimes referred to as a "discrete gradient vector field" in the literature; in particular, this is the case in [10].

2.2. **Discrete Morse Theory and Homology.** The discrete Morse theory sketched in Subsection 2.1 can be used to compute homology groups. We include only a brief account, without proofs. More extended expositions can be found in [13] and [12].

Fix an oriented finite regular CW complex X. Let $C_*(X)$ be the cellular chain complex of X. Each chain group $C_n(X)$ has a distinguished basis consisting of positively oriented n-cells, denoted $B_n(X)$. Let W be a Morse matching, and define a map $\widehat{W}_n : C_n(X) \to C_{n+1}(X)$ as follows:

$$\begin{aligned} \widehat{W}_n(c) &= \pm W(c) \text{ if } c \text{ is redundant;} \\ \widehat{W}_n(c) &= 0 \text{ otherwise.} \end{aligned}$$

Here the sign is chosen so that the oriented cell c occurs with the coefficient -1 in $\partial \widehat{W}_n(c)$ if c is redundant. Extend linearly to a map $\widehat{W}_n : C_n(X) \to C_{n+1}(X)$. Define a chain map $f_{\widehat{W}} : C_*(X) \to C_*(X)$, called the *discrete flow* associated to $\widehat{W}$, by setting $f_{\widehat{W}} = 1 + \partial \widehat{W} + \widehat{W} \partial$. We usually omit the subscript and simply write f.

The discrete flow f has the following properties:

Lemma 2.1.

(1) *([12], [13]) For any finite chain $c \in C_*(X)$, there is some $m \in \mathbb{N}$ such that $f^m(c) = f^{m+1}(c) = \ldots$. It follows that there is a well-defined chain map $f^\infty : C_*(X) \to C_*(X)$.*

(2) ([12]; cf. [13]) *If c is any cycle in $C_*(X)$, then there is a unique f-invariant cycle that is homologous to c, namely $f^\infty(c)$. Moreover, $f^\infty(c)$ is a linear combination of oriented critical cells and collapsible cells (i.e., any redundant cell appears with a coefficient of 0).*

(3) ([12]) *If c is a collapsible cell, then $f^\infty(c) = 0$. If c is critical, then $f^\infty(c) = c + (collapsible\ cells)$. As a result, an f-invariant chain is determined by its critical cells, i.e., if c is an f-invariant chain and $c = c_{crit} + c_{coll}$, where c_{crit} is a linear combination of critical cells and c_{coll} is a linear combination of collapsible cells, then*

$$c = f^\infty(c) = f^\infty(c_{crit}).$$

Properties (1)-(3) show that if a finite regular CW complex X is endowed with a Morse matching W, then the homology groups of X are largely determined by the critical cells of X. We now make this statement more precise. Fix a Morse matching W. For $i \geq 0$, let $M_i(X)$ denote the free abelian group on the set of positively oriented critical i-cells. Give the collection of abelian groups $M_i(X)$ ($i \geq 0$) the structure of a chain complex, called the *Morse complex*, by identifying $M_i(X)$ with $C_i(X)$ via the map f^∞. The boundary map $\tilde{\partial}$ in the Morse complex is defined by

$$\tilde{\partial}(c) = \Pi \partial f^\infty(c) \quad (c \in M_i(X)),$$

where Π denotes projection onto the factor of $C_{i-1}(X)$ spanned by the critical $i-1$ cells.

We have the following theorem:

Theorem 2.2. ([12], [13]) *The Morse complex $\left(M_n(X), \tilde{\partial}_n \right)$ and the cellular chain complex $(C_n(X), \partial_n)$ have isomorphic homology groups, by an isomorphism which sends a cycle c from the Morse complex to $f^\infty(c)$.* □

3. A Morse Matching on the Discretized Configuration Space $UD^n\Gamma$

3.1. **Definitions and an Example.** Throughout this paper, all graphs are assumed to be finite and connected.

Let Γ be a graph, and fix a natural number n. The *labelled configuration space* of Γ on n points is the space

$$\left(\prod^n \Gamma \right) - \Delta,$$

where Δ is the set of all points $(x_1, \ldots, x_n) \in \prod^n \Gamma$ such that $x_i = x_j$ for some $i \neq j$. The *unlabelled configuration space* of Γ on n points is the quotient of the labelled configuration space by the action of the symmetric group S_n, where the action permutes the factors. The *braid group* of Γ on n strands, denoted $B_n\Gamma$, is the fundamental group of the unlabelled configuration space of Γ on n strands. The *pure braid group*, denoted $PB_n\Gamma$, is the fundamental group of the labelled configuration space.

The set of vertices of Γ will be denoted by $V(\Gamma)$, and the degree of a vertex $v \in V(\Gamma)$ is denoted $d(v)$. If a vertex v is such that $d(v) \geq 3$, v is called *essential*.

Let Δ' denote the union of those open cells of $\prod^n \Gamma$ whose closures intersect Δ. Let $\mathcal{D}^n\Gamma$ denote the space $\prod^n \Gamma - \Delta'$. Note that $\mathcal{D}^n\Gamma$ inherits a CW complex structure from the Cartesian product, and that a cell in $\mathcal{D}^n\Gamma$ has the form $c_1 \times \cdots \times c_n$ such that each c_i is either a vertex or the interior of an edge, and the

closures of the c_i are mutually disjoint. Let $U\mathcal{D}^n\Gamma$ denote the quotient of $\mathcal{D}^n\Gamma$ by the action of the symmetric group S_n which permutes the coordinates. Thus, an open cell in $U\mathcal{D}^n\Gamma$ has the form $\{c_1, \ldots, c_n\}$ such that each c_i is either a vertex or the interior of an edge and the closures are mutually disjoint. The set notation is used to indicate that order does not matter.

Under most circumstances, the labelled (respectively, unlabelled) configuration space of Γ is homotopy equivalent to $\mathcal{D}^n\Gamma$ (respectively, $U\mathcal{D}^n\Gamma$). Specifically:

Theorem 3.1. [1] *For any $n > 1$ and any graph Γ with at least n vertices, the labelled (unlabelled) configuration space of n points on Γ strong deformation retracts onto $\mathcal{D}^n\Gamma$ ($U\mathcal{D}^n\Gamma$) if*

(1) *each path between distinct vertices of degree not equal to 2 passes through at least $n-1$ edges; and*

(2) *each path from a vertex to itself which is not null-homotopic in Γ passes through at least $n+1$ edges.*

A graph Γ satisfying the conditions of this theorem for a given n is called *sufficiently subdivided* for this n. It is clear that every graph is homeomorphic to a sufficiently subdivided graph, no matter what n may be.

Throughout the rest of the paper, we work exclusively with the space $U\mathcal{D}^n\Gamma$ where Γ is sufficiently subdivided for n. Also from now on, "edge" and "cell" will refer to closed objects.

Choose a maximal tree T in Γ. Edges outside of T are called *deleted edges*. Pick a vertex $*$ of valence 1 in T to be the root of T. Choose an embedding of the tree T into the plane. We define an order on the vertices of T (and, thus, on vertices of Γ) as follows. Begin at the basepoint $*$ and walk along the tree, following the leftmost branch at any given intersection, and consecutively number the vertices in the order in which they are first encountered. (When you reach a vertex of degree one, turn around.) The vertex adjacent to $*$ is assigned the number 1. Note that this numbering depends only on the choice of $*$ and the embedding of the tree. Let $\iota(e)$ and $\tau(e)$ denote the endpoints of a given edge e of Γ. Without loss of generality, we orient each edge to go from $\iota(e)$ to $\tau(e)$, and so that $\iota(e) > \tau(e)$. (Thus, if $e \subseteq T$ the geodesic segment $[\iota(e), *]$ in T must pass through $\tau(e)$.)

We use the order on the vertices to define a Morse matching W on $U\mathcal{D}^n\Gamma$. We begin with some definitions which will help to classify cells of $U\mathcal{D}^n\Gamma$ as critical, collapsible, or redundant.

Let $c = \{c_1, \ldots, c_{n-1}, v\}$ be a cell in $U\mathcal{D}^n\Gamma$ containing a vertex v. If $v = *$, then v is *blocked* in c; otherwise, let e be the unique edge in T such that $\iota(e) = v$. If $e \cap c_i \neq \emptyset$ for some $i \in \{1, \ldots, n-1\}$, we also say v is *blocked* in c; otherwise, v is unblocked. Equivalently, v is unblocked in c if and only if $\{c_1, \ldots, c_{n-1}, e\}$ is also a cell in $U\mathcal{D}^n\Gamma$. If $c = \{c_1, \ldots, c_{n-1}, e\}$, the edge e is *disrespectful* in c if

(1) there is a vertex v in c such that
   (a) v is adjacent to $\tau(e)$, and
   (b) $\tau(e) < v < \iota(e)$, or
(2) e is a deleted edge.

Otherwise, the edge e is *respectful* in c. Conceptually, think of the edge e as representing a strand in c moving from $\iota(e)$ to $\tau(e)$. Then e is disrespectful in c if that strand is moving out of turn by not respecting the order on vertices of T. In the paper [10], disrespectful was referred to by "non order-respecting".

It will occasionally be useful to have another definition. If v is a vertex in the tree T, we say that two vertices v_1 and v_2 lie in the same *direction* from v if the geodesics $[v, v_1], [v, v_2] \subseteq T$ start with the same edge. Thus, there are $deg(v)$ directions from a vertex of degree $deg(v)$ in T. We number these directions $0, 1, 2, \ldots, deg(v) - 1$, beginning with the direction represented by $[v, *]$, numbered 0, and proceeding in clockwise order. We will sometimes write $g(v_1, v_2)$ (where $v_1 \neq v_2$) to refer to the direction from v_1 to v_2.

Suppose that we are given a cell $c = \{c_1, \ldots, c_n\}$ in $UD^n\Gamma$. Assign each cell in c a number as follows. A vertex of c is given the number from the above traversal of T. An edge e of c is given the number for $\iota(e)$. Arrange the cells of c in a sequence $\mathcal{S}$, from the least- to the greatest-numbered. The following definition of a Morse matching W is equivalent to the definition of W from [10], by Theorem 3.6 of the same paper.

Definition 3.2. We define a Morse matching W on $U\mathcal{D}^n\Gamma$ as follows:
   (1) If an unblocked vertex occurs in $\mathcal{S}$ before all of the respectful edges in c (if any), then $W(c)$ is obtained from c by replacing the minimal unblocked vertex $v \in c$ with $e(v)$, where $e(v)$ is the unique edge in T satisfying $\iota(e(v)) = v$. In particular, c is redundant.
   (2) If a respectful edge occurs before any unblocked vertex, then $c \in \text{im } W$, i.e., c is collapsible. The cell $W^{-1}(c)$ is obtained from c by replacing the minimal respectful edge e with $\iota(e)$.
   (3) If there are neither unblocked vertices nor respectful edges in c, then c is critical.

Example 3.3. Figure 1 depicts three different cells of UD^4T_{min} for the given tree T_{min}. In each case, the vertices and edges of the given cell are numbered from least to greatest, in the sense mentioned above. (The numbering of these cells differs from the above-described order, but this doesn't matter since the ordering remains the same. For instance, the vertices and edges in the cell pictured in Figure 1 (a) should be numbered 10, 14, 16, and 19, instead of (respectively) 1, 2, 3, 4.)


FIGURE 1. Three different cells of $U\mathcal{D}^nT$.

The vertex numbered 1 in (a) is blocked. The vertex numbered 2 is unblocked, so the cell in (a) is redundant. Note that edge 3 is respectful, and edge 4 is disrespectful. We get $W(c_1)$ by replacing vertex 2 with the unique edge in T having vertex 2 as its initial vertex. In terms of the usual ordering, this is the edge $[v_{14}, v_{13}]$.

Let c_2 denote the cell depicted in (b). The vertex numbered 1 in c_2 is blocked. The edge numbered 2 is respectful, so c_2 is collapsible. Note that vertex 3 is blocked and edge 4 is disrespectful. The description of W^{-1} above implies that $W^{-1}(c_2)$ is obtained from c_2 by replacing edge 2 with its initial vertex.

The cell depicted in (c) is critical since vertices 1 and 2 are both blocked, and the edges 3 and 4 are disrespectful.

4. The Mod 2 Cohomology Ring of $U\mathcal{D}^n T$

In this section, we give a partial description of the mod 2 cohomology ring of $U\mathcal{D}^n T$, where n is an arbitrary natural number and T is an arbitrary tree. The method is to map $U\mathcal{D}^n T$ to a new complex $\widehat{U\mathcal{D}^n T}$, which is a subcomplex of a high-dimensional torus. The induced map $q^* : H^*(\widehat{U\mathcal{D}^n T}) \to H^*(U\mathcal{D}^n T)$ turns out to be surjective. This gives us an easy way to compute the cup product: we take two cohomology classes in $H^*(U\mathcal{D}^n T)$, look at their preimages under the map q^*, cup these preimages using known facts about the cohomology rings of subcomplexes of torii, and then push the product back over into $H^*(U\mathcal{D}^n T)$.

The complex $\widehat{U\mathcal{D}^n T}$ can be described very simply: it is the result of identifying the opposite sides of all of the cubes in $U\mathcal{D}^n T$. Thus $\widehat{U\mathcal{D}^n T}$ consists of a union of (potentially singular) torii, one for each cell of $U\mathcal{D}^n T$.

It is by no means clear, however, and false in general, that identifying all opposite faces in a CAT(0) cubical complex will result in a subcomplex of a torus. For this reason, we give a very careful (and somewhat abstract) proof that $\widehat{U\mathcal{D}^n T}$ has the properties we want.

4.1. An Equivalence Relation on the Cells of $U\mathcal{D}^n T$.
Let $E(c)$ denote the set of edges of the i-cell c. We abuse the notation by also letting $E(c)$ denote the subset $\bigcup_{e \in E(c)} e$ of T.

Let c and c' be i-cells ($0 \le i \le n$) of $U\mathcal{D}^n T$. Say $c = \{e_1, \dots, e_i, v_{i+1}, \dots, v_n\}$ and $c' = \{e'_1, \dots, e'_i, v'_{i+1}, \dots, v'_n\}$, where $E(c) = \{e_1, \dots, e_i\}$ and $E(c') = \{e'_1, \dots, e'_i\}$. Write $c \sim > c'$ if

(1) $E(c) = E(c')$, and
(2) for any connected component C of $T - E(c)$,
$$\left| C \cap \{v_{i+1}, \dots, v_n\} \right| = \left| C \cap \{v'_{i+1}, \dots, v'_n\} \right|.$$

Let K be the set of open cells in $U\mathcal{D}^n T$, as in Section 2.1. It is rather clear that $\sim$ is an equivalence relation on K. Let $[c]$ denote the equivalence class of a cell c.

We define a partial order $\le$ on the equivalence classes based on the partial order on cells, writing $[c] \le [c_1]$ if there exist representatives $\hat{c} \in [c]$, $\hat{c}_1 \in [c_1]$ such that $\hat{c} \le \hat{c}_1$. It is slightly nontrivial to verify that $\le$ is transitive on the equivalence classes. Suppose that $[c_1] \le [c_2]$ and $[c_2] \le [c_3]$. There are representatives $\hat{c}_1 \in [c_1], \hat{c}_2, \tilde{c}_2 \in [c_2]$, and $\tilde{c}_3 \in [c_3]$ such that $\hat{c}_1 \le \hat{c}_2$ and $\tilde{c}_2 \le \tilde{c}_3$. Since $\hat{c}_2 \sim \tilde{c}_2$, it is possible, by moving the vertices of $\hat{c}_2$ one at a time along edges of T, and leaving all edges fixed, to arrive at $\tilde{c}_2$. Every vertex of $\hat{c}_2$ is also a vertex of $\hat{c}_1$; if a vertex of $\hat{c}_1$ is also in $\hat{c}_2$, then move it as above. Call the result of doing these moves $\tilde{c}_1$. Then $\tilde{c}_1 \sim \hat{c}_1$ and it is clear that $\tilde{c}_1 \le \tilde{c}_2 \le \tilde{c}_3$, so transitivity of $\le$ follows.

We state the following lemma in terms of a tree T, but we note that with the appropriate definitions the statements in parts (1), (4), and (5) may be generalized to arbitrary graphs.

Lemma 4.1 (Properties of $\sim$ and $\leq$). *Let T be a finite connected tree.*

  (1) *If c and c' are i-cells of $U\mathcal{D}^n T$ with $E(c) = E(c')$ and the equivalence classes $[c], [c']$ have a common upper bound $[\tilde{c}]$ with respect to $\leq$, then $[c] = [c']$.*
  (2) *Let $c_1, \ldots, c_j$ be 1-cells in $U\mathcal{D}^n T$ from distinct equivalence classes. If $\{[c_1], \ldots, [c_j]\}$ has an upper bound with respect to $\leq$, then $\{[c_1], \ldots, [c_j]\}$ has a least upper bound with respect to $\leq$. Furthermore, if $e_1, \ldots, e_j$ are edges of T satisfying $e_i \in c_i$ for $1 \leq i \leq j$, then $e_1, \ldots, e_j$ are pairwise disjoint.*
  (3) *If $\tilde{c}$ is a j-cell in $U\mathcal{D}^n T$, then there is a unique collection $\{[c_1], \ldots, [c_j]\}$ of equivalence classes of 1-cells such that $[\tilde{c}]$ is the least upper bound of $\{[c_1], \ldots, [c_j]\}$ with respect to $\leq$.*
  (4) *If c is a critical cell in $U\mathcal{D}^n T$ and $[c'] \leq [c]$, then $c' \sim \hat{c}$ for some critical cell $\hat{c}$.*
  (5) *If c is a critical cell in $U\mathcal{D}^n T$ and $c' \sim c$, then $c' = c$ or c' is redundant.*

Proof. (1) Let $[\tilde{c}]$ be an equivalence class of j-dimensional cells in $U\mathcal{D}^n T$, where $\tilde{c} = \{e_1, \ldots, e_j, v_{j+1}, \ldots, v_n\}$. It is enough to show that, for any c with $[c] \leq [\tilde{c}]$, the equivalence class $[c]$ is uniquely determined by the collection $\{e_{i_1}, \ldots, e_{i_k}\}$ of all edges common to both $\tilde{c}$ and c.

An arbitrary face c of $\tilde{c}$ is determined by selecting a subset of the edges $\{e_1, \ldots, e_j\}$ and replacing each edge of this subset by either its initial or its terminal vertex. By an argument similar to that establishing the transitivity of $\leq$, the equivalence class of c depends neither on the representative chosen from $[\tilde{c}]$ nor on the choice involved in replacing an edge with one of its endpoints. Thus, given $\tilde{c}$ with $[c] \leq [\tilde{c}]$, $[c]$ is uniquely determined by the edges c has in common with $\tilde{c}$ - i.e. $[c]$ is uniquely determined by $E(c)$. This proves part (1).

(2) Suppose that $[c]$ is an upper bound for $\{[c_1], \ldots, [c_j]\}$ where the $[c_1], \ldots, [c_j]$ are all distinct. For $i \in \{1, \ldots, j\}$, let e_i be the unique edge in c_i. Certainly, $\{e_1, \ldots, e_j\} \subseteq c$. Let c' be a cell given by replacing any extra edges $e \in c - \{e_1, \ldots, e_j\}$ with either $\iota(e)$ or $\tau(e)$. Then $[c_i] \leq [c']$ for $i \in \{1, \ldots, j\}$, so $[c']$ is an upper bound for $\{[c_1], \ldots, [c_j]\}$. It follows that given any upper bound $[c]$ for $\{[c_1], \ldots, [c_j]\}$ there exists another upper bound $[c']$ such that $[c'] \leq [c]$ and $E(c') = \{e_1, \ldots, e_j\}$.

To prove the first claim of part (2), it remains to be shown that there is only one upper bound $[c']$ such that $E(c') = \{e_1, \ldots, e_j\}$. Suppose that $[c']$ and $[c'']$ are both such upper bounds for $\{[c_1], \ldots, [c_j]\}$. Thus $E(c') = E(c'') = \{e_1, \ldots, e_j\}$. Fix an integer i, $1 \leq i \leq j$. Let C be a connected component of $T - e_i$, and assign to C the integer $f_{e_i}(C)$ defined by:

$$f_{e_i}(C) \quad = \quad (\#\{\text{vertices or edges of } c' \text{ contained in } C\})$$
$$- (\#\{\text{vertices or edges of } c'' \text{ contained in } C\}).$$

It must be that $f_{e_i}(C) = 0$ for every i and every component C of $T - e_i$. For if not, then for some i there exist distinct equivalence classes $[c'_0], [c''_0]$ of 1-cells with $E(c'_0) = E(c''_0) = \{e_i\}$ such that $[c'_0] \leq [c']$ and $[c''_0] \leq [c'']$. But by part (1), we must have that $[c_i] = [c'_0]$ since $[c']$ is a common upper bound for $[c_i]$ and $[c'_0]$. Similarly, $[c_i] = [c''_0]$. This means that $[c'_0] = [c''_0]$, which is a contradiction, so indeed $f_{e_i}(C) = 0$ for any i and any C.

If $[c'] \neq [c'']$, then c' and c'' have a different number of cells in some connected component of $T - E(c')$. Fix a connected component C of $T - E(c')$. Let $e_{i_1}, \ldots, e_{i_l}$

be the edges of the collection $\{e_1, \ldots, e_j\}$ that have exactly one endpoint in the component C. We analyze the connected components of $T - \bigcup_{k=1}^l e_{i_k}$. These connected components are either C itself, or contain a single endpoint from exactly one of the edges $e_{i_1}, \ldots, e_{i_l}$. (Note: here we have just used the fact that T is a tree for the first time.) If a connected component C' of $T - \bigcup_{k=1}^l e_{i_k}$ is not C itself, and contains an endpoint of e_{i_m}, say, then it contains equal numbers of cells from c' and c'' by the claim, since C' is in fact a connected component of $T - e_{i_m}$. (This again uses the fact that T is a tree.) It follows by process of elimination that C contains equal numbers of cells, necessarily vertices, from both c' and c''. Since C was an arbitrary connected component of $T - E(c')$, it must be that $[c'] = [c'']$.

We now prove the second claim. Suppose that $[c_1], \ldots, [c_j]$ are distinct equivalence classes of 1-cells having a common upper bound $[c]$. Part (1) shows that $e_{i_1} \neq e_{i_2}$ for two distinct equivalence classes $[c_{i_1}], [c_{i_2}] \in \{[c_1], \ldots, [c_j]\}$. If $e_{i_1} \in c_{i_1}$, $e_{i_2} \in c_{i_2}$, $e_{i_1} \cap e_{i_2} \neq \emptyset$, then $e_{i_1}, e_{i_2} \in c$, which is impossible since c is a cell of $U\mathcal{D}^n T$ and $e_{i_1} \cap e_{i_2} \neq \emptyset$.

(3) Let $S = \{[c] \mid \dim c = 1 \text{ and } [c] \leq [\tilde{c}]\}$. Part (1) implies that an element of S is uniquely determined by a choice of edge from $\tilde{c}$. Thus $|S| = j$. The fact that $[\tilde{c}]$ is the least upper bound of S follows from the description of the least upper bound in (2).

(4) If c is a critical cell of $U\mathcal{D}^n \Gamma$ and $[c_1] < [c]$, then there is some representative of the equivalence class $[c_1]$ - say c_1 - such that $c_1 < c$. For, let $c_1' \in [c_1]$ and $c' \in [c]$ be such that $c_1' < c'$. Then c_1 (respectively, c) is the result of moving all vertices in c_1' (respectively, c') toward $*$ until they are blocked. Thus each edge in c_1 is an edge in c, and each vertex in c is a vertex in c_1. It follows that no edges in c_1 are respectful. Now repeatedly move each unblocked vertex of c_1 toward $*$ until it is blocked. This operation clearly preserves $\sim$, and the resulting cell is critical, having no unblocked vertices and no respectful edges.

(5) Suppose c is critical and $c_1 \sim c$. Suppose first that c_1 has no respectful edges. If c_1 has unblocked vertices, then it follows that c_1 is redundant. If c_1 has no unblocked vertices, then it is critical by Definition 3.2. In fact $c_1 = c$ in this case, since both cells involve the same edges, the vertices in both are blocked, and each component of $T - E(c)$ contains the same number of vertices from each of c_1 and c.

Now suppose that c_1 has respectful edges. Let e be the smallest such (recall that "smallest" means that $\iota(e)$ is minimal). Since the edge e is disrespectful in c, there is some vertex $v \in c$ adjacent to $\tau(e)$ and satisfying

$$0 < g(\tau(e), v) < g(\tau(e), \iota(e)).$$

Let C be the connected component of $T - E(c)$ containing $\tau(e)$ and lying in the direction $g(\tau(e), v)$ from $\tau(e)$. This component contains vertices of c and thus vertices of c_1, since $c_1 \sim c$. If C contains unblocked vertices of c_1, then any such vertex v_1 satisfies $v_1 < \iota(e)$, and so it follows that c_1 is redundant. If c contains only blocked vertices, then it follows that the vertex v is a vertex of c_1, whence the edge e is disrespectful in c_1, a contradiction. $\square$

4.2. The complex $\widehat{U\mathcal{D}^n T}$.

Define a complex $\widehat{U\mathcal{D}^n T}$ as follows. For each equivalence class $[c]$ of 1-cells in the set K of open cells of $U\mathcal{D}^n T$, introduce a copy of S^1, denoted $S^1_{[c]}$. Give $S^1_{[c]}$ a cell structure with one open 1-cell, denoted $e^1_{[c]}$, and

one 0-cell. Form the finite product $\prod_{[c]} S^1_{[c]}$. Since we are interested in giving each cell an explicit characteristic map, we order the factors in the product as follows. Assign to each equivalence class $[c]$ of 1-cells the number $N([c])$ of the vertex $\iota(e)$, where e is the unique edge satisfying $e \in c$. This numbering of equivalence classes is well-defined (though not one-to-one). Arrange the factors of $\prod_{[c]} S^1_{[c]}$ so that if $N([c_1]) < N([c_2])$, then the factor $S^1_{[c_1]}$ occurs before $S^1_{[c_2]}$. (This arrangement of factors is not unique.)

Since each 1-cell of this product corresponds naturally to an equivalence class $[c]$ of 1-cells in $U\mathcal{D}^n T$, each i-cell corresponds to a collection $\{[c_1], \ldots, [c_i]\}$ of (distinct) equivalence classes of 1-cells. We obtain the space $\widehat{U\mathcal{D}^n T}$ by throwing out an open i-cell $\{[c_1], \ldots, [c_i]\}$ if $\{[c_1], \ldots, [c_i]\}$ has no upper bound. If $\{[c_1], \ldots, [c_i]\}$ has an upper bound, then it has a least upper bound $[c]$, and we label the corresponding i-cell by $[c]$. Note that, by Lemma 4.1(3), the equivalence classes $[c]$ are in one-to-one correspondence with cells of $\widehat{U\mathcal{D}^n T}$, and the dimension of the cell c of $U\mathcal{D}^n T$ is the same as that of the cell labelled $[c]$ in $\widehat{U\mathcal{D}^n T}$.

If R is a field, then the *exterior algebra* ([16], pg. 217) on a set $\{v_1, v_2, \ldots, v_{n+1}\}$, denoted $\Lambda_R[v_1, \ldots, v_{n+1}]$, is the R-module having relations $v_i v_j = -v_j v_i$ and $v_i^2 = 0$. The products $v_{i_1} v_{i_2} \ldots v_{i_j}$ ($0 \le j \le n$; $i_1 < i_2 < \ldots < i_j$) form a basis. The empty product is the multiplicative identity.

For any equivalence class $[c]$ of j-cells, let $\hat{\phi}_{[c]} : C_*(U\hat{\mathcal{D}}^n T) \to \mathbb{Z}/2\mathbb{Z}$ denote the j-cocycle satisfying $\hat{\phi}_{[c]}([c]) = 1$ and $\hat{\phi}_{[c]}([c']) = 0$ for all $[c'] \ne [c]$. (Here $C_j(U\hat{\mathcal{D}}^n T)$ denotes the free abelian group generated by oriented j-cells of $U\hat{\mathcal{D}}^n T$. Note that $\hat{\phi}_{[c]}$ is necessarily a cocycle, since all of the boundary maps in the cellular chain complex of a subcomplex of $\prod S^1$ are 0.)

Proposition 4.2. *The space $\widehat{U\mathcal{D}^n T}$ is a CW complex. The mod 2 cohomology ring $H^*\left(\widehat{U\mathcal{D}^n T}; \mathbb{Z}/2\mathbb{Z}\right)$ is isomorphic to*

$$\Lambda_{\mathbb{Z}/2\mathbb{Z}}\big[[c_1], \ldots, [c_k]\big]/I,$$

where $\{[c_1], \ldots, [c_k]\}$ is the collection of all equivalence classes of 1-cells in $U\mathcal{D}^n T$, and I is the ideal generated by the set of all monomial terms $[c_{i_1}] \cdot [c_{i_2}] \cdot \ldots \cdot [c_{i_m}]$ where $\{[c_{i_1}], \ldots, [c_{i_m}]\}$ has no upper bound.

The isomorphism sends the j-cocycle $\hat{\phi}_{[c]}$ to the product $[c_{i_1}] \cdot [c_{i_2}] \cdot \ldots \cdot [c_{i_j}]$, where $\{[c_{i_1}], \ldots, [c_{i_j}]\}$ is the unique collection with least upper bound $[c]$ as in Lemma 4.1(3).

The elements $\hat{\phi}_{[c]}$ form a basis for the cohomology as $[c]$ ranges over all possible equivalence classes.

Proof. The complex $\widehat{U\mathcal{D}^n T}$ inherits a cell structure from $\prod_{[c]} S^1_{[c]}$. To prove that $\widehat{U\mathcal{D}^n T}$ is a CW complex, we need to verify that, for every open cell in $\widehat{U\mathcal{D}^n T}$, the attaching map to $\prod_{[c]} S^1_{[c]}$ also maps into $\widehat{U\mathcal{D}^n T}$. This means we must show that if an open cell $\{[c_{j_1}], \ldots, [c_{j_l}]\}$ is thrown out of the product $\prod_{[c]} S^1_{[c]}$, then so is any other open cell having $\{[c_{j_1}], \ldots, [c_{j_l}]\}$ as a face. By definition, if the cell $\{[c_{j_1}], \ldots, [c_{j_l}]\}$ is thrown out, the collection $\{[c_{j_1}], \ldots, [c_{j_l}]\}$ has no upper bound. Any cell having $\{[c_{j_1}], \ldots, [c_{j_l}]\}$ as a face must be labelled by a collection S of equivalence classes of 1-cells satisfying $\{[c_{j_1}], \ldots, [c_{j_l}]\} \subseteq S$. But as $\{[c_{j_1}], \ldots, [c_{j_l}]\}$ has no upper bound, S can have no upper bound. Since S has no upper bound, by

the defininition of $\widehat{U\mathcal{D}^n T}$ the open cell labelled S is also thrown out of the product, as required.

The remaining statements then follow easily from the description of $\widehat{U\mathcal{D}^n T}$ as a subcomplex of $\prod_{[c]} S^1_{[c]}$, and from the description in [16] (pg. 227) of the cohomology rings of subcomplexes of the torus.   □

4.3. **A map** $q : U\mathcal{D}^n T \to \widehat{U\mathcal{D}^n T}$ **and the induced map on cohomology.** For each edge e in the tree T, choose a characteristic map $h_e : [0,1] \to e$, such that $h_e(0) = \iota(e)$ and $h_e(1) = \tau(e)$. These maps induce characteristic maps on the cells of $U\mathcal{D}^n T$ as follows. Let $c = \{e_1, \ldots, e_i, v_{i+1}, \ldots, v_n\}$ be an i-cell of $U\mathcal{D}^n T$, and suppose without loss of generality that $\iota(e_1) < \iota(e_2) < \ldots < \iota(e_i)$. Define the characteristic map $h_c : [0,1]^i \to c$ by

$$h_c(t_1, \ldots, t_i) = \{h_{e_1}(t_1), \ldots, h_{e_i}(t_i), v_{i+1}, \ldots, v_n\}.$$

Note that $h_c(t_1, \ldots, t_i)$ is an n-element subset of T here, rather than an n-tuple of cells in T. The map h_c is a homeomorphism.

We now choose characteristic maps for the cells of $\widehat{U\mathcal{D}^n T}$. Begin by choosing a characteristic map $h_{[c]} : [0,1] \to S^1_{[c]}$ for each 1-cell $e^1_{[c]}$. Suppose that $[c]$ is the label of an i-dimensional cell in $\widehat{U\mathcal{D}^n T}$. Thus, $[c]$ is the least upper bound of a collection $\{[c_1], \ldots, [c_i]\}$ of distinct equivalence classes of 1-cells. If $e_1, \ldots, e_i$ are the unique edges satisfying $e_j \in c_j$ $(1 \le j \le i)$, then $e_1, \ldots, e_i$ are pairwise disjoint by Lemma 4.1 (2). In particular, the natural numbers $N([c_1]), \ldots, N([c_i])$ are all different. We assume, without loss of generality, that $N([c_1]) < \ldots < N([c_i])$. The characteristic map for the cell $[c]$ in $\widehat{U\mathcal{D}^n T}$ is $\hat{h}_{[c]} : [0,1]^i \to e^i_{[c]}$, defined by $\hat{h}_{[c]}(t_1, \ldots, t_i) = (h_{[c_1]}(t_1), \ldots, h_{[c_i]}(t_i))$, where the value of $\hat{h}_{[c]}(t_1, \ldots, t_i)$ on each omitted factor $[\hat{c}]$, where $[\hat{c}]$ is an equivalence class of 1-cells, is always assumed to be the unique vertex of $S^1_{[\hat{c}]}$.

Now we are ready to define the map $q : U\mathcal{D}^n T \to \widehat{U\mathcal{D}^n T}$. Consider the following diagram:

FIGURE 2. Defining the map q.

The vertical arrow is a quotient map, so, by a well-known principle (e.g., [7], Theorem 3.2), there will exist a well-defined map q making the above diagram commute if $\coprod_{c \in K} \hat{h}_{[c]}$ is constant on point inverses of $\coprod_{c \in K} h_c$.

Proposition 4.3. *There is a well-defined map* $q : U\mathcal{D}^n T \to \widehat{U\mathcal{D}^n T}$ *making the above diagram commute. The map* $q^* : H^* \left(\widehat{U\mathcal{D}^n T}; \mathbb{Z}/2\mathbb{Z} \right) \to H^* \left(U\mathcal{D}^n T; \mathbb{Z}/2\mathbb{Z} \right)$ *sends* $\hat{\phi}_{[c]}$ *to the cohomology class of* $\phi_{[c]} \in C^* \left(U\mathcal{D}^n T; \mathbb{Z}/2\mathbb{Z} \right)$, *where* $\phi_{[c]}$ *is the*

cellular cocycle satisfying:

$$\phi_{[c]}(\tilde{c}) = 1 \qquad if \ \tilde{c} \sim c$$
$$\phi_{[c]}(\tilde{c}) = 0 \qquad otherwise \ .$$

Proof. We have to show that $\coprod_{c \in K} \hat{h}_{[c]}$ is constant on point inverses of $\coprod_{c \in K} h_c$. For this, it is sufficient to show that $\hat{h}_{[c']} \circ h_{c'}^{-1} = \hat{h}_{[c]} \circ h_c^{-1} \mid_{c'}$, where c' is a codimension-one face of c. We let c be a j-dimensional cell in $U\mathcal{D}^n T$, say $c = \{e_1, \ldots, e_j, v_{j+1}, \ldots, v_n\}$, where the edges e_i are arranged in order. Assume, without loss of generality, that $c' = \{\iota(e_1), e_2, \ldots, e_j, v_{j+1}, \ldots, v_n\}$.

Choose a point $x \in c'$. (Here we really mean a point in the cell c', rather than one of the "members" $\iota(e_1), e_2, \ldots, e_j, v_{j+1}, \ldots, v_n$ of c'.) Suppose that $x = h_{c'}(t_1, \ldots, t_{j-1})$, i.e., $x = \{\iota(e_1), h_{e_2}(t_1), \ldots, h_{e_j}(t_{j-1}), v_{j+1}, \ldots, v_n\}$.

$$
\begin{aligned}
\hat{h}_{[c']} \circ h_{c'}^{-1}(x) &= \hat{h}_{[c']}(t_1, \ldots, t_{j-1}) \\
&= \left(h_{[c'_1]}(t_1), \ldots, h_{[c'_{j-1}]}(t_{j-1}) \right).
\end{aligned}
$$

$$
\begin{aligned}
\hat{h}_{[c]} \circ h_c^{-1}(x) &= \hat{h}_{[c]}(0, t_1, \ldots, t_{j-1}) \\
&= \left(h_{[c_1]}(0), h_{[c_2]}(t_1), \ldots, h_{[c_j]}(t_{j-1}) \right) \\
&= \left(h_{[c_2]}(t_1), \ldots, h_{[c_j]}(t_{j-1}) \right).
\end{aligned}
$$

For this last equality, recall that $h_{[c_1]}(0)$ is the vertex of $S^1_{[c_1]}$, and we omit such factors for the sake of simplicity.

Now $[c'_i]$, by definition, is the unique equivalence class of 1-cells satisfying: (i) $e_{i+1} \in c'_i$, and (ii) $[c'_i] \leq [c]$. Note that $[c_{i+1}]$ has the same properties, so $[c'_i] = [c_{i+1}]$. It follows that the map q exists.

The remaining statements about cohomology follow easily from the description of the map $q : U\mathcal{D}^n T \to \widehat{U\mathcal{D}^n T}$. The main point is that the interior of a cell c in $U\mathcal{D}^n T$ is mapped homeomorphically to the interior of $[c]$, and thus the mod 2 mapping number of c with $[c]$ is equal to 1. $\qquad\square$

To describe the mod 2 cohomology ring of $U\mathcal{D}^n T$, we will need to recall a result from [11]:

Theorem 4.4. ([11], *Theorem 3.7*) *The boundary maps in the Morse complex* $(M_*(U\mathcal{D}^n T), \tilde{\partial})$ *are all zero. In particular, $H_i(U\mathcal{D}^n T)$ is a free abelian group of rank equal to the number of critical i-cells in $U\mathcal{D}^n T$.* $\qquad\square$

By Lemma 2.1, we obtain explicit cycles in $C_i(U\mathcal{D}^n T)$, the ith cellular chain group of $U\mathcal{D}^n T$, by applying the map f^∞ to any linear combination of critical i-cells. A collection of representatives for a distinguished basis of the cellular i-dimensional homology is thus $\{f^\infty(c) \mid c \ is \ a \ critical \ i-cell\}$. For the sake of simplicity in notation, we express a cellular homology class in $H_*(U\mathcal{D}^n T)$ as a linear combination $\Sigma a_i c_i$ of critical cells c_i, as opposed to $f^\infty(\Sigma a_i c_i)$. Identify $H^i(U\mathcal{D}^n T; \mathbb{Z}/2\mathbb{Z})$ with $Hom(H_i(U\mathcal{D}^n T); \mathbb{Z}/2\mathbb{Z})$ by the universal coefficient isomorphism. Let c^* denote the dual of a critical cell c.

Proposition 4.5.

(1) If c is a critical cell in $U\mathcal{D}^n T$, then $q^*(\hat{\phi}_{[c]}) = c^*$.

(2) Let c be a critical cell in $U\mathcal{D}^n T$. If $[c]$ is the least upper bound of $\{[c_1], \ldots, [c_i]\}$, where the $[c_1], \ldots, [c_i]$ are distinct equivalence classes of 1-cells, then, without loss of generality, $c_1, \ldots, c_i$ are critical and

$$c_1^* \cup \ldots \cup c_i^* = c^*.$$

(3) If $[c_1], \ldots, [c_i]$ are distinct equivalence classes of 1-cells having the least upper bound $[c]$, then

$$\left[\phi_{[c_1]}\right] \cup \ldots \cup \left[\phi_{[c_i]}\right] = \left[\phi_{[c]}\right].$$

If $[c_1], \ldots, [c_i]$ have no upper bound, then the above cup product is 0.

Proof. (1) By Proposition 4.3, $\phi_{[c]}$ is a cocycle representative of $q^*(\hat{\phi}_{[c]})$, where $\phi_{[c]}(\tilde{c}) = 1$ if $\tilde{c} \sim c$, and $\phi_{[c]}(\tilde{c}) = 0$ otherwise. Note that, by Lemma 4.1(5), the support of $\phi_{[c]}$ consists of redundant cells, and a single critical cell (c itself), but no collapsible cells.

We evaluate the cohomology class of $\phi_{[c]}$ on a basis for $H_i(U\mathcal{D}^n T)$ consisting of critical cells $c_1, \ldots, c_j$. As c is the unique critical cell in the support of $\phi_{[c]}$, for a critical 1-cell c_k (viewed as a homology class), $\phi_{[c]}(c_k) = 1$ if and only if $c_k = c$. The statement of (1) follows.

(2) The statement that $c_1, \ldots, c_i$ may be chosen to be critical follows from Lemma 4.1(4). By Proposition 4.2, $\hat{\phi}_{[c]} = \hat{\phi}_{[c_1]} \cup \ldots \cup \hat{\phi}_{[c_i]}$ in $H^*(\widehat{U\mathcal{D}^n T}; \mathbb{Z}/2\mathbb{Z})$. By (1), the statement of (2) follows.

(3) This follows from applying Propositions 4.2 and 4.3. $\qquad\square$

4.4. **A computation of** $H^*(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z})$. Let T_{min} be the tree depicted in Figure 3. The tree T_{min} is the tree, unique up to homeomorphism, with the fewest number of essential vertices which is not 'linear': i.e., the vertices of degree 3 or more in T_{min} do not all lie on a single embedded line segment. We compute the mod 2 cohomology ring of $U\mathcal{D}^4 T_{min}$ as an application of the ideas of this section. The results will be used in the proof of Theorem 5.11.

To begin, we will need to compute the integral homology groups of $U\mathcal{D}^4 T_{min}$. According to Proposition 4.1 of [11], we have

$$H_0(U\mathcal{D}^4 T_{min}) \cong \mathbb{Z} \quad , \quad H_1(U\mathcal{D}^4 T_{min}) \cong \mathbb{Z}^{24}$$
$$H_2(U\mathcal{D}^4 T_{min}) \cong \mathbb{Z}^6 \quad , \quad H_n(U\mathcal{D}^4 T_{min}) \cong 0 (n \geq 3).$$

We also need to describe the critical cells c, which are determined by the following choices. First, choose the locations of the edges of c. The requirement that the edges in c be disrespectful means that there are only four possibilities: $[v_3, v_7], [v_9, v_{19}], [v_{12}, v_{16}]$, and $[v_{21}, v_{25}]$. We let e_k denote the unique edge in T having v_k as its initial vertex. With this notation, we can rewrite the four edges above as e_7, e_{19}, e_{16}, and e_{25}, respectively. A choice of n edges from this collection, together with the requirement that the edges be disrespectful, forces n vertices also to be in c. More specifically, $v_4 \in c$ if $e_7 \in c$, $v_{10} \in c$ if $e_{19} \in c$, $v_{13} \in c$ if $e_{16} \in c$, and $v_{22} \in c$ if $e_{25} \in c$. It immediately follows that there are exactly $\binom{4}{2} = 6$ critical 2-cells and no critical n-cells for $n \geq 3$, so $H_2(U\mathcal{D}^4 T_{min}) \cong \mathbb{Z}^6$ and $H_3(U\mathcal{D}^4 T_{min}) \cong 0$ if $n \geq 3$.

If c is a critical 1-cell, there is another choice to make. We have determined the location of an edge e and a vertex v. The locations of the other two vertices are completely determined by specifying how many are in each component of $T -$

FIGURE 3. The minimal nonlinear tree T_{min}.

$\{\tau(e)\}$, since all vertices in c are blocked. There are 3 distinguishable components of $T - \{\tau(e)\}$ and 2 remaining indistinguishable vertices, which make 6 possible ways. There are thus a total of 24 critical 1-cells c, since there are 4 possible choices for the edge e in c, 6 choices for the remaining vertices and these choices are independent.

Finally, we note that there is only one critical 0-cell.

To understand the multiplication in $H^*(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z})$, it is clearly enough to understand the product of any two elements c_1^*, c_2^* (c_1 and c_2 are critical 1-cells) of the standard dual basis for $H^1(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z})$. Proposition 4.5(1) says that $c_1^* = [\phi_{[c_1]}]$ and $c_2^* = [\phi_{[c_2]}]$, and Proposition 4.5(3) says that $[\phi_{[c_1]}] \cup [\phi_{[c_2]}] = 0$ if $c_1 \sim c_2$, or if $\{[c_1], [c_2]\}$ has no upper bound. We are therefore led to determine which equivalence classes $[c]$ of 2-cells can be the upper bound for a pair of distinct equivalence classes of 1-cells $[c_1]$ and $[c_2]$, where c_1 and c_2 are critical.

For this, it will be helpful to have a definition. If c is a j-cell in $U\mathcal{D}^n\Gamma$, and c' is obtained from c by replacing each member of some collection $\{e_{i_1}, \ldots, e_{i_m}\} \subseteq E(c)$ of edges with either its initial or its terminal vertex, then c' is the result of *breaking* the edges $\{e_{i_1}, \ldots, e_{i_m}\}$ in c. We note that the choice of replacing a given edge e_{i_l} with $\tau(e_{i_l})$ or $\iota(e_{i_l})$ is made independently for each edge, and these choices do not affect $[c']$. In fact, by Lemma 4.1 (1), $[c']$ is completely determined by its edges, since $[c'] \leq [c]$. Note also that c' is a codimension-m face of c, and conversely if c' is a codimension-m face of c then c' is obtained from c by breaking m edges.

Suppose that $[c]$ is the least upper bound for two distinct equivalence classes $[c_1]$ and $[c_2]$, where c_1 and c_2 are critical 1-cells in $U\mathcal{D}^4 T_{min}$. For $i = 1, 2$, let e_i be the unique edge in c_i. (Note: this is inconsistent with the convention that $\iota(e_k) = v_k$, but should cause no confusion.) By Lemma 4.1(2), e_1 and e_2 are disjoint edges, and, by the above description of critical cells in $U\mathcal{D}^4 T_{min}$, $\{e_1, e_2\} \subseteq \{e_7, e_{19}, e_{16}, e_{25}\}$. Thus, if $[c]$ is an upper bound for $[c_1]$ and $[c_2]$ where c_1 and c_2 critical, we have:

FIGURE 4. A preliminary picture of the relations in $H^*(U\mathcal{D}^n T_{min}; \mathbb{Z}/2\mathbb{Z})$. The numbers on the edges refer to the numbering of 2-cells in the text. The vertices represent duals of critical 1-cells. Here, B denotes the critical 1-cell $\{e_{19}, v_{12}, v_{11}, v_{10}\}$, B' denotes the critical 1-cell $\{e_{19}, v_{11}, v_{10}, *\}$, and C denotes the critical 1-cell $\{e_{16}, v_{13}, v_1, *\}$.

(i) the edges of c, namely e_1 and e_2, are distinct elements of the above 4-element set, and (ii) the cell resulting from breaking either of the edges e_1, e_2 in c must be equivalent to a critical cell. Reformulating slightly, we get the following conditions, which must be satisfied by c:

(1) $\{e_1, e_2\} \subseteq \{e_7, e_{19}, e_{16}, e_{25}\}$.
(2) If $e_7 \in c$, then v_4, v_5, or v_6 is in c.
(3) If $e_{16} \in c$, then v_{13}, v_{14}, or v_{15} is in c.
(4) If $e_{25} \in c$, then v_{22}, v_{23}, or v_{24} is in c.
(5) If $e_{19} \in c$, then either $e_{16} \in c$ or at least one element of $\{v_{10}, \ldots, v_{18}\}$ is in c.

In light of (3), (5) may be replaced with:

(5') If $e_{19} \in c$, then at least one element of $\{v_{10}, \ldots, v_{18}\}$ is in c.

A total of 10 distinct equivalence classes of 2-cells have representatives c satisfying (1)-(5). Representatives of these classes are:

(1) $\{e_7, e_{19}, v_4, v_{10}\}$		(6) $\{e_{19}, e_{16}, v_{13}, v_{14}\}$
(2) $\{e_7, e_{16}, v_4, v_{13}\}$		(7) $\{e_{19}, e_{16}, v_{13}, v_{17}\}$
(3) $\{e_7, e_{25}, v_4, v_{22}\}$		(8) $\{e_{19}, e_{16}, v_{13}, v_{10}\}$
(4) $\{e_{16}, e_{25}, v_{13}, v_{22}\}$		(9) $\{e_{19}, e_{16}, v_{13}, v_{20}\}$
(5) $\{e_{19}, e_{25}, v_{10}, v_{22}\}$		(10) $\{e_{19}, e_{16}, v_{13}, *\}$

The reason is that a choice of edges e_1, e_2 from $\{e_7, e_{16}, e_{19}, e_{25}\}$ completely determines $[c]$, by (1)-(5), unless $\{e_1, e_2\} = \{e_{16}, e_{19}\}$. The first five edges listed above result from the five cases in which $\{e_1, e_2\} \neq \{e_{16}, e_{19}\}$. If $\{e_1, e_2\} = \{e_{16}, e_{19}\}$,

then one of the vertices of c must be v_{13}, v_{14} or v_{15} (by (3)), but the other vertex may be chosen from any of the five remaining components of $T_{min} - (e_{16} \cup e_{19})$, and this accounts for the last five 2-cells above.

Figure 4 below depicts these 10 equivalence classes $[c]$ of 2-cells as line segments whose endpoints are the equivalence classes $[c_1], [c_2]$ of 1-cells (c_1, c_2 critical) satisfying $[c_1], [c_2] \leq [c]$.

It follows from what we've said so far that there are only 10 equivalence classes $[c']$ of 1-cells such that: (1) $[c']$ contains a critical 1-cell, and (2) there is another distinct equivalence class $[c'']$, also containing a critical 1-cell, such that $\{[c'], [c'']\}$ has an upper bound. By Proposition 4.5(1) & (3), these equivalence classes correspond to the only elements of the standard basis for $H^1(U\mathcal{D}^4T_{min}; \mathbb{Z}/2\mathbb{Z})$ which might have non-trivial cup products. In fact, checking labels of edges, it is not difficult to see that cases (1)-(5) and (8) are all (distinct) critical cells, and thus correspond to linearly independent elements of the standard basis for $H^2(U\mathcal{D}^4T_{min}; \mathbb{Z}/2\mathbb{Z})$. We now consider the remaining edges.

Lemma 4.6. *If $\hat{c}$ is one of the 2-cells (6), (7) or (9), then $\phi_{[\hat{c}]}$ represents 0 in cohomology.*

Proof. Consider the following cochains $\alpha_6, \alpha_7, \alpha_9 : C_1(U\mathcal{D}^4T_{min}) \to \mathbb{Z}/2\mathbb{Z}$, depicted pictorially in Figure 5. The cochain α_6 sends a given 1-cell c to 1 if and only if: (1)


(a) α_6 (b) α_7 (c) α_9

(d) α_{10}

FIGURE 5. The 1-cochains α_6, α_7, α_9, and $\alpha_1 0$ used in Lemmas 4.6 and 4.7. Each 1-cochain maps a given 1-cell c to 1 in $\mathbb{Z}/2\mathbb{Z}$ if and only if c contains the edge shown and has exactly as many strands specified in each circled portion of the tree.

c contains the edge e_{16}; (2) c contains exactly one element from $\{*, v_1, \ldots, v_{11}\}$, and (3) c contains exactly two elements from $\{v_{13}, v_{14}, v_{15}\}$. (Of course, α_6 sends any other 1-cell to 0.) The cochain α_7 sends a given 1-cell c to 1 if and only if c satisfies conditions (1) and (2) in the definition of α_6, as well as: (3') c contains exactly one element from $\{v_{17}, v_{18}\}$, and exactly one element from $\{v_{13}, v_{14}, v_{15}\}$. The cochain α_9 sends a given 1-cell c to 1 if and only if c contains exactly two elements from the set $\{v_{19}, \ldots, v_{27}\}$, exactly one element from the set $\{v_{13}, v_{14}, v_{15}\}$, and the edge e_{16}.

We leave it as an exercise to show that the coboundaries $\delta(\alpha_6)$, $\delta(\alpha_7)$, and $\delta(\alpha_9)$, are precisely $\phi_{[\hat{c}]}$, where $\hat{c}$ is the cell (6), (7), and (9), respectively. $\qquad\square$

Lemma 4.7. *Let c_1 be the cell labelled (8), and let c_2 be the cell labelled (10). The cocycles $\phi_{[c_1]}$ and $\phi_{[c_2]}$ are cohomologous in $H^1(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z})$.*

Proof. Consider the 1-cochain α_{10} defined as follows and depicted in Figure 5. Let α_{10} be the 1-cochain which sends a 1-cell c to 1 if and only if: (1) c contains the edge e_{16}; (2) c contains exactly one of the vertices $\{v_{13}, v_{14}, v_{15}\}$, and (3) c contains exactly two vertices from $\{*, v_1, \ldots, v_{11}\}$.

We leave it as an exercise to show that $\delta(\alpha_{10}) = \phi_{[c_1]} + \phi_{[c_2]}$. $\qquad\square$

We now interpret Figure 4 as a multiplication table for the cup product. If we let dashed edges correspond to 0 products and perform an elementary row operation, we arrive at Figure 6.


FIGURE 6. A picture of the cohomology ring $H^*(U\mathcal{D}^n T_{min}; \mathbb{Z}/2\mathbb{Z})$. The solid edges represent duals of critical 2-cells; the vertices represent (sums of) duals of critical 1-cells. Two vertices cup to the solid edge connecting them, or to 0 if there is no such edge.

We thus arrive at a complete description of the multiplication in $H^*(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z})$: an 18-dimensional subspace W of $H^1(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z})$ annihilates all one-dimensional cohomology classes. This subspace W is spanned by

the duals of the 14 critical 1-cells not appearing in Figure 6, together with the four elements of $H^1 (U\mathcal{D}^n T; \mathbb{Z}/2\mathbb{Z})$ which are endpoints of only dashed lines. The multiplication in the remaining 6-dimensional subspace is described by the subgraph of Figure 6 consisting of six solid lines and the six vertices they connect: two basis elements cup to the label of the solid edge connecting them, or to 0 if there is no such edge.

This description of the multiplication in $H^* \left(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z} \right)$ suggests that it is an exterior face algebra, an idea we define in the next section.

5. Counterexamples to Ghrist's Conjecture

5.1. **Preliminaries on Exterior Face Rings.** For a ring R, an *exterior ring* over R on a set $\{v_1, v_2, \dots, v_n\}$, denoted $\Lambda_R[v_1, \dots, v_n]$, is the free R-module having the products $v_{i_1} v_{i_2} \dots v_{i_j}$ $(0 \leq i_1 < i_2 < \dots < i_j \leq n)$ as a basis. The empty product is the multiplicative identity. The multiplicative relations are generated by all relations of the following types: $v_i v_j = -v_j v_i$ and $v_i^2 = 0$.

Let $K = (\{v_1, \dots, v_n\}, S)$ be a finite simplicial complex and R be a commutative ring with identity. The *exterior face ring* $\Lambda_R (K)$ of K over R is the quotient of the exterior ring $\Lambda_R[v_1, \dots, v_n]$ by the relations $v_{i_1} \dots v_{i_k} = 0$ for $0 \leq i_1 < i_2 < \dots < i_k \leq n$ and when $\{v_{i_1}, \dots, v_{i_k}\} \notin S$. Note that an exterior ring on a set $\{v_1, v_2, \dots, v_n, v_{n+1}\}$ is isomorphic the exterior face ring $\Lambda_R (K)$ where K is a standard n-simplex. If R is a field, then $\Lambda_R(K)$ inherits an algebra structure, and is called an *exterior face algebra*.

Example 5.1. The calculation of the previous subsection shows that the ring $H^*(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z})$ is isomorphic to $\Lambda_{\mathbb{Z}/2\mathbb{Z}}(K)$, where K is the union of 18 isolated vertices with a graph isomorphic to the one in Figure 6 (2) consisting of the six solid edges labelled (1) - (5) and (8), and vertices incident with them.

If $R = \mathbb{Z}/2\mathbb{Z}$, then the exterior face ring $\Lambda_{\mathbb{Z}/2\mathbb{Z}} (K)$ is a quotient of a polynomial ring:
$$\Lambda_{\mathbb{Z}/2\mathbb{Z}} (K) = \mathbb{Z}/2\mathbb{Z} [v_1, \dots, v_n] / I(K),$$
where $I(K)$ is the ideal of $\mathbb{Z}/2\mathbb{Z} [v_1, \dots, v_n]$ generated by the set
$$\left\{ v_1^2, \dots, v_n^2 \right\} \cup \left\{ v_{i_1} v_{i_2} \dots v_{i_k} \mid \{v_{i_1}, \dots, v_{i_k}\} \notin S \right\}.$$
In this case, since R is a field, we have that $\Lambda_R(K)$ inherits an algebra structure. Throughout the rest of the paper, all exterior face rings $\Lambda_R(K)$ will be over $\mathbb{Z}/2\mathbb{Z}$, and we will therefore drop the subscript R without further comment.

A simplicial complex K is *flag* if, whenever a collection of vertices $v_{i_1}, \dots, v_{i_j} \in K$ pairwise span edges, $\{v_{i_1}, \dots, v_{i_j}\}$ is a simplex of K.

In case K is a flag complex, there is a simple set of generators for $I(K)$:

Lemma 5.2. *If K is a flag complex, then*
$$I(K) = \left\langle \left\{ v_1^2, \dots, v_n^2 \right\} \cup \left\{ v_{i_1} v_{i_2} \mid \{v_{i_1}, v_{i_2}\} \notin S \right\} \right\rangle.$$

Proof. Let $I'(K)$ denote the ideal on the right half of the equality in the lemma. We need to show that $I(K) \subseteq I'(K)$, the reverse inclusion being obvious.

Suppose that $v_{i_1} \dots v_{i_k}$ satisfies $\{v_{i_1}, \dots, v_{i_k}\} \notin S$. Since K is a flag complex, there must exist $v_{j_1}, v_{j_2} \in \{v_{i_1}, \dots, v_{i_k}\}$ such that $\{v_{j_1}, v_{j_2}\} \notin S$. It follows that $v_{j_1} v_{j_2} \in I'(K)$. Now $v_{j_1} v_{j_2} \mid v_{i_1} \dots v_{i_k}$, so $v_{i_1} \dots v_{i_k} \in I'(K)$. Thus, $I(K) \subseteq I'(K)$. $\qquad\square$

Example 5.3. Let Γ be a finite simple graph. The *right-angled Artin group* G_Γ associated to Γ is a group defined by a presentation in which the generators are in one-to-one correspondence with vertices of Γ, and relations consist of all commutators of the form $[v_i, v_j]$, where v_i and v_j are adjacent in Γ.

Charney and Davis [4] have described $K(G_\Gamma, 1)$ complexes for all right-angled Artin groups (generalizing the Salvetti complex of [19] for spherical Artin groups). Begin with a torus $\prod S^1$, where the factors are in one-to-one correspondence with vertices in Γ. Assume that each S^1 is given the standard cellulation, consisting of one 0-cell and one 1-cell. Their $K(G_\Gamma, 1)$ space is obtained from this product by throwing out an open i-cell if the i 1-cells in its factorization correspond to vertices $v_1, v_2, \ldots, v_i$ which do not form a clique, i.e., if some pair of vertices $v_{j_1}, v_{j_2} \in \{v_1, \ldots, v_i\}$ do not span an edge of Γ.

This description of $K(G_\Gamma, 1)$, together with the description of the cohomology rings of subcomplexes of a torus in [16] (pg. 227), implies

Proposition 5.4. *The cohomology ring* $H^*(G_\Gamma; \mathbb{Z}/2\mathbb{Z})$ *is the exterior face ring* $\Lambda(K)$, *where* K *is the unique flag complex having* Γ *as its 1-skeleton.* $\square$

(Here K is the simplicial complex whose n-simplices are the cliques in Γ having $n+1$ members.)

We can now give a simple principle which will allow us to find counterexamples to Ghrist's Conjecture 1.1. A homomorphism $\phi : R[x_1, \ldots, x_l] \to R[y_1, \ldots, y_m]$ between polynomial rings is *degree-preserving* if it sends any homogeneous polynomial of degree k to another homogeneous polynomial of degree k (or, equivalently, if it sends any homogeneous polynomial of degree 1 to another homogeneous polynomial of degree 1). More generally, if R_1 and R_2 are quotients of polynomial rings by ideals generated by homogeneous polynomials, then $\phi : R_1 \to R_2$ is *degree-preserving* if any equivalence class of homogeneous polynomials is mapped to an equivalence class of homogeneous polynomials of the same degree.

Proposition 5.5. *Let* K *be a flag complex, and let* $\partial\Delta^n$ *be the boundary of the standard n-simplex ($n \geq 2$). If* $\phi : \Lambda(K) \to \Lambda(\partial\Delta^n)$ *is a degree-preserving surjection, then* $\ker\phi$ *cannot be generated by homogeneous degree 1 and degree 2 elements.*

Proof. Let $\{v_1, \ldots, v_m\}$ be the vertices of K. The hypotheses imply that ϕ induces a linear surjection from the space $\Lambda(K)^1$ of homogeneous degree 1 elements of $\Lambda(K)$ to the space $\Lambda(\partial\Delta^n)^1$ of homogeneous degree 1 elements of $\Lambda(\partial\Delta^n)$, which is $(n+1)$−dimensional. Thus there is a collection of $n+1$ elements of the standard basis $\{v_1, \ldots, v_m\}$ for $\Lambda(K)$ which map onto a basis for the space of homogeneous degree 1 elements of $\Lambda(\partial\Delta^n)$. We can thus assume, without loss of generality, that $\{\phi(v_1), \ldots, \phi(v_{n+1})\}$ is a basis for $\Lambda(\partial\Delta^n)^1$.

For $i \in \{n+2, \ldots, m\}$, let s_i denote the (unique) linear combination of $v_1, \ldots, v_{n+1}$ such that $\phi(s_i) = \phi(v_i)$. Thus, each $s_i + v_i$ is an element of $\ker\phi$. Since the set $\{s_i + v_i \mid i \in \{n+2, \ldots, m\}\}$ is linearly independent, it must form a basis for $\ker\phi \cap \Lambda(K)^1$, since the dimension of $\ker\phi \cap \Lambda(K)^1$ is $m - n - 1$.

Now assume that $\ker\phi$ is generated by degree 1 and degree 2 elements. Suppose that $\ker\phi \cap \Lambda(K)^2$ is spanned by $t_1, t_2, \ldots, t_k$, where t_i, for $i \in \{1, \ldots, k\}$, is a homogeneous element of degree 2. By Lemma 5.2,

$$\Lambda(K) \cong \Lambda[v_1, \ldots, v_m] / \langle u_1, \ldots, u_l \rangle,$$

where u_i is a homogeneous element of degree 2 for $1 \leq i \leq l$. It follows that

$$\phi : \Lambda\left[v_1, \ldots, v_m\right] / \left\langle s_{n+2} + v_{n+2}, \ldots, s_m + v_m, t_1, \ldots t_k, u_1, \ldots, u_l\right\rangle \to \Lambda\left(\partial \Delta^n\right)$$

is an isomorphism. Let $\Lambda_\phi(K)$ be the quotient of $\Lambda\left[v_1, \ldots, v_{n+1}\right]$ by the ideal $I_\phi(K) = \langle \hat{t}_1, \ldots, \hat{t}_k, \hat{u}_1, \ldots, \hat{u}_l \rangle$, where $\hat{t}_j$ (respectively, $\hat{u}_j$) is the result of replacing v_i with s_i $(n + 2 \leq i \leq m)$ in t_j (respectively, u_j). Note that $I_\phi(K)$ is generated by homogeneous elements of degree 2. It is easy to see that the map

$$\psi : \Lambda_\phi(K) \to \Lambda\left[v_1, \ldots, v_m\right] / \left\langle s_{n+2} + v_{n+2}, \ldots, s_m + v_m, t_1, \ldots t_k, u_1, \ldots, u_l\right\rangle,$$

sending v_i to v_i for $i \in \{1, \ldots, n + 1\}$, is an isomorphism, which also preserves degree.

Now we obtain a contradiction by counting the dimensions of $\Lambda_\phi(K)^2$, $\Lambda_\phi(K)^{n+1}$, $\Lambda\left(\partial \Delta^n\right)^2$, and $\Lambda\left(\partial \Delta^n\right)^{n+1}$ as vector spaces. We have:

$$\dim\left(\Lambda\left(\partial \Delta^n\right)^2\right) = \frac{n(n+1)}{2}; \qquad \dim\left(\Lambda\left(\partial \Delta^n\right)^{n+1}\right) = 0.$$

Either $I_\phi(K)$ is the 0 ideal or it isn't. If it is, then

$$\dim\left(\Lambda_\phi(K)^{n+1}\right) = 1;$$

if it isn't, then

$$\dim\left(\Lambda_\phi(K)^2\right) < \frac{n(n+1)}{2}.$$

In either case, we have a contradiction since $\phi \circ \psi$ is a degree-preserving bijection and thus preserves the dimension in each degree. $\square$

Corollary 5.6. *Let K_1 and K_2 be finite simplicial complexes.*
   (1) If $\phi : \Lambda(K_1) \to \Lambda(K_2)$ is a degree-preserving surjection, K_1 is a flag complex, and $\ker \phi$ is generated by homogeneous elements of degrees one and two, then K_2 is also a flag complex.
   (2) If $\phi : \Lambda(K_1) \to \Lambda(K_2)$ is a degree-preserving isomorphism, then K_1 is a flag complex if and only if K_2 is.

Proof. (1) If K_2 is not flag, then for some $n \geq 2$, $\partial \Delta^n$ is a full subcomplex of K_2, i.e., $\partial \Delta^n$ is not the boundary of an n-simplex in K_2. Define a map $\psi : \Lambda(K_2) \to \Lambda(\partial \Delta^n)$, sending a given vertex v to 0 if $v \notin \partial \Delta^n$, and to itself otherwise. The map ψ is a degree-preserving surjection whose kernel is generated by elements of degree 1. It follows that $\psi \circ \phi : \Lambda(K_1) \to \Lambda(\partial \Delta^n)$ is a degree-preserving surjection whose kernel is generated by homogeneous elements of degrees 1 and 2. This contradicts Proposition 5.5.

(2) This is an easy consequence of (1). $\square$

Note that Gubeladze [15] has proven a strong generalization of Corollary 5.6(2): $\phi : \Lambda(K_1) \to \Lambda(K_2)$ is a degree-preserving isomorphism if and only if K_1 is isomorphic to K_2 as simplicial complexes.

Our experience in computing $H^*\left(U\mathcal{D}^n T; \mathbb{Z}/2\mathbb{Z}\right)$ for various small examples, including the case $T = T_{min}$ and $n = 4$, suggests the following conjecture:

Conjecture 5.7. *The cohomology ring $H^*\left(U\mathcal{D}^n T; \mathbb{Z}/2\mathbb{Z}\right)$ is an exterior face algebra, for any tree T and any n.*

We note finally that the conjecture seems just as likely to be true for arbitrary fields, not simply $\mathbb{Z}/2\mathbb{Z}$.

5.2. **Which Tree Braid Groups Are Right-Angled Artin?** Recall the definition of a right-angled Artin group from Example 5.3. In this subsection, we characterize exactly which tree braid groups are right-angled Artin. Theorem 5.11 states that a tree braid group $B_n T$ is a right-angled Artin group exactly when either $n < 4$ or T is linear (recall that a tree is *linear* if there exists an embedded line segment which contains all of the essential vertices of T).

Let T_{min} be the minimal nonlinear tree described in Subsection 4.4.

Lemma 5.8. *Let $n \geq 4$. Let T be a nonlinear tree that is sufficiently subdivided for n. There is a cellular embedding θ of (a suitably subdivided) T_{min} into T such that:*

(1) *the image of T_{min} is sufficiently subdivided for 4 strands;*
(2) *there is a choice of basepoints $\overline{\ast}$ for T_{min} and $\ast$ for T such that $\overline{\ast}$ has degree 1 in T_{min}, $\ast$ has degree 1 in T, and the geodesic segment $[\ast, \overline{\ast}]$ in T crosses exactly $n - 4$ edges, none of which are edges of T_{min}.*

Proof. Choose a collection $\mathcal{C}$ of essential vertices of T such that the elements of $\mathcal{C}$ all lie along an embedded arc, and such that $\mathcal{C}$ is a maximal set of essential vertices with this property. Fix an arc $[v_1, v_2]$ in T such that $\mathcal{C} \subseteq [v_1, v_2]$, where v_1 and v_2 are essential. Since T is a nonlinear tree, there exists an essential vertex $v_3 \notin [v_1, v_2]$. Consider the geodesic segment γ connecting v_3 to $[v_1, v_2]$. By maximality of $\mathcal{C}$, γ must meet $[v_1, v_2]$ in another essential vertex $v_4 \notin \{v_1, v_2\}$, for otherwise $\gamma \cup [v_1, v_2]$ is an arc containing $\mathcal{C} \cup \{v_3\}$, which contradicts the maximality of $\mathcal{C}$.


FIGURE 7. The larger tree T is sufficiently subdivided for $n = 6$. The smaller tree (encircled) is a copy of T_{min}. This figure shows the image of a critical 2-cell in $U\mathcal{D}^4 T_{min}$ under the map $\theta : U\mathcal{D}^4 T_{min} \to U\mathcal{D}^6 T$.

The Y-graph formed by the segments $[v_1, v_4]$, $[v_2, v_4]$, and $[v_3, v_4]$ is sufficiently subdivided for 4, since the tree T is sufficiently subdivided for n and $n \geq 4$. For $i = 1, 2, 3$, add to the Y-graph two additional embedded line segments at v_i, each consisting of exactly 3 edges, in such a way that the new segments have no edges in common with either each other or with the Y-graph. It is possible to do this

because each of the vertices v_1, v_2, and v_3 are essential. The result of this procedure gives a cellular embedding of T_{min} into T which satisfies (1).

To produce an embedding satisfying (2) as well, proceed as follows. Choose a vertex $\hat{*}$ having degree 1 in T_{min}. If $\hat{*}$ has degree 1 in T, then since T is sufficiently subdivided for n it must be that $n = 4$. In this case, choose $* = \overline{*} = \hat{*}$. Otherwise, $\hat{*}$ has degree at least 2 in T. Choose an arc $\hat{\gamma}$ in T with no edges in common with the embedding of T_{min}, and with the embedding of $\hat{*}$ as one of its endpoints. Furthermore, choose $\hat{\gamma}$ to be a maximal such arc, so that the other endpoint of $\hat{\gamma}$ has degree 1 in T. Declare the other endpoint of $\hat{\gamma}$ to be $*$, and let $\overline{*}$ be the (unique) vertex lying on $\hat{\gamma}$ at distance exactly $n - 4$ from $*$. Modify T_{min} (if necessary) by adding in the segment $[\overline{*}, \hat{*}]$.   □

Let $n \geq 4$, and let T be a nonlinear tree. The embedding θ of the previous lemma induces a map of configuration spaces $\theta : U\mathcal{D}^4 T_{min} \rightarrow U\mathcal{D}^n T$, defined by $\theta(\{c_1, c_2, c_3, c_4\}) = \{\theta(c_1), \theta(c_2), \theta(c_3), \theta(c_4)\} \cup \{*, v_1, v_2, \ldots, v_{n-5}\}$, where $*, v_1, v_2, \ldots, v_{n-5}$ are the $n - 4$ vertices of T closest to $*$ (see Figure 7).

Choose an embedding of T in the plane, and consider the induced classifications of cells in $U\mathcal{D}^n T$ and $U\mathcal{D}^4 T_{min}$ into the critical, collapsible, and redundant cell types.

Proposition 5.9. *The map θ preserves cell type - i.e. takes critical, collapsible, and redundant cells in $U\mathcal{D}^4 T_{min}$ to critical, collapsible, and redundant cells in $U\mathcal{D}^n T$, respectively.*

Proof. Let c be a cell in $U\mathcal{D}^4 T_{min}$. Since the map θ on configuration spaces is induced by a cellular embedding on the level of trees, by the choices of $*$ and $\overline{*}$, a vertex in c is blocked if and only if it is blocked in $\theta(c)$. Since the embedding of T_{min} in the plane is induced from the embedding of T in the plane, an edge in c is respectful if and only if it is respectful in $\theta(c)$. The cell $\theta(c)$ is obtained from c by adding exactly $n - 4$ blocked vertices. Thus, the numbering on vertices and edges in $\theta(c)$ used to determine cell type (see the discussion preceding Definition 3.2) differs from the numbering for c only by the insertion of $n - 4$ blocked vertices at the beginning of the numbering. By the definition of the Morse matching, Definition 3.2, the proposition is proven.   □

Proposition 5.10. *We have:*
  (1) The map θ induces an injection $\theta_* : H_* (U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z}) \rightarrow H_* (U\mathcal{D}^n T; \mathbb{Z}/2\mathbb{Z})$. The homology class corresponding to a given critical cell $c \subseteq U\mathcal{D}^4 T_{min}$ goes to a homology class corresponding to $\theta(c) \subseteq U\mathcal{D}^n T$. In particular, the image of θ_* is a direct factor of $H_* (U\mathcal{D}^n T)$.
  (2) The induced map $\theta^* : H^*(U\mathcal{D}^n T; \mathbb{Z}/2\mathbb{Z}) \rightarrow H^*(U\mathcal{D}^4 T_{min}; \mathbb{Z}/2\mathbb{Z})$ sends the dual of a critical cell c to $(\theta^{-1}(c))^*$ if c is in the image of θ, or to 0 otherwise.

Proof. (1) Fix an embedding of T into the plane, and choose an embedding $\theta : T_{min} \rightarrow T$ as in Lemma 5.8. We note that, due to the choices of the embedding $\theta : T_{min} \rightarrow T$ and basepoints, the map $\theta : U\mathcal{D}^4 T_{min} \rightarrow U\mathcal{D}^n T$ sends collapsible cells to collapsible cells, redundant cells to redundant cells, and critical cells to critical cells.

If c is an arbitrary critical cell of $U\mathcal{D}^4 T_{min}$, then a cycle representing the homology class determined by c is $f^\infty(c)$, which has the form $c + (collapsible\ cells)$. Since θ preserves a cell's type, it follows that the homology class $\theta_*(c)$ may be represented by

a cycle of the form $\theta(c)+(collapsible\ cells)$, where $\theta(c)$ is critical. By Lemma 2.1(2), the cycle $\theta(c) + (collapsible\ cells)$ is homologous to $f^\infty(\theta(c) + (collapsible\ cells))$. By Lemma 2.1(3) and the fact that $\theta(c)$ is critical,

$$f^\infty(\theta(c) + collapsible\ cells) = \theta(c) + ((different)\ collapsible\ cells).$$

On the other hand, the homology class corresponding to the critical cell $\theta(c)$ is, by definition, $f^\infty(\theta(c))$, which consists of $\theta(c)+(collapsible\ cells)$. Thus, using the fact from Lemma 2.1(3) that an f-invariant chain is determined by its critical cells, we conclude that $\theta_*(c) = \theta(c)$, as required.

(2) This is an easy consequence of (1) and the naturality of the universal coefficient isomorphism. $\qquad\square$

Theorem 5.11. *The tree braid group B_nT is a right-angled Artin group if and only if T is linear or $n < 4$.*

Proof. ($\Leftarrow$) Connolly and Doig [5] showed that B_nT is a right-angled Artin group if T is linear. If $n < 4$ and T is a tree, then Theorem 4.3 of [10] shows that B_nT is in fact a free group, since $U\mathcal{D}^nT$ strong deformation retracts on a graph. This proves one direction.

(Note that it is possible to get a proof of Connolly and Doig's result as an application of the ideas in [10]. Suppose T is a linear tree. Choose some basepoint $*$ for T and an embedded arc ℓ containing $*$ and all essential vertices of T; it is possible to do this since T is linear. Now embed T in $\mathbb{R}^2$ so that: (1) $*$ is mapped to the origin; (2) ℓ is mapped to a segment on the positive y-axis, and (3) the image of T is contained in $\{(x,y) \in \mathbb{R}^2 \mid x \leq 0\}$. With this choice of embedding and the induced order on the vertices of T, Theorem 5.3 of [10] gives a presentation of B_nT as a right-angled Artin group. The proof is left as an exercise for the interested reader.)

($\Rightarrow$) Proof by contradiction. Suppose that T is nonlinear, $n \geq 4$, and B_nT is a right-angled Artin group. Since $U\mathcal{D}^nT$ is aspherical ([1], [14]), $U\mathcal{D}^nT$ is a $K(B_nT, 1)$. In particular, by Proposition 5.4, the cohomology ring $H^*(U\mathcal{D}^nT; \mathbb{Z}/2\mathbb{Z})$ is the exterior face algebra of a flag complex.

We choose an embedding $\theta : U\mathcal{D}^4T_{min} \to U\mathcal{D}^nT$ as in Lemma 5.8. By Proposition 5.10, $\theta^* : H^*(U\mathcal{D}^nT; \mathbb{Z}/2\mathbb{Z}) \to H^*(U\mathcal{D}^4T_{min}; \mathbb{Z}/2\mathbb{Z})$ is surjective, and it is necessarily degree-preserving. Since $H^*(U\mathcal{D}^4T_{min}; \mathbb{Z}/2\mathbb{Z})$ is the exterior face algebra of a complex that is not flag, we will arrive at a contradiction to Corollary 5.6(1) if we can show that $\ker(\theta^*)$ is generated by homogeneous elements of degrees one and two. For this, it is sufficient to show that if c is a critical cell in $U\mathcal{D}^nT$ of dimension at least 3, then c^* is divisible by some element $c_1^* \in \ker(\theta^*)$ of degree one.

Let c be a critical cell in $U\mathcal{D}^nT$ of dimension at least 3. There are two cases: either every cell of c lies inside of (the embedded image of) $T_{min} \cup [*, \overline{*}]$, or some cell of c is not contained in $T_{min} \cup [*, \overline{*}]$.

We first consider the case in which some vertex or edge x of T occurring in c is not contained in $T_{min} \cup [*, \overline{*}]$. Either x is an edge e or x is a blocked vertex. If x is a blocked vertex, then at the largest essential vertex on the geodesic $[x, *]$ there must be a disrespectful edge e. In either case, break all edges of c other than e, and consider the resulting 1-cell c'. By Lemma 4.1(4), c' is equivalent to a critical 1-cell $\tilde{c}'$, and the proof of Lemma 4.1(4) shows that $\tilde{c}'$ may be described as simply the result of moving all vertices in c' toward $*$ until they are all blocked. If follows that

x occurs in $\tilde{c}'$. This implies that $\tilde{c}'$ is not in the image of $\theta : U\mathcal{D}^4 T_{min} \to U\mathcal{D}^n T$, since all cells in this image consist of cells in T_{min}. It now follows from Proposition 5.10(2) that $\tilde{c}'^* \in \ker(\theta^*)$. But $[c]$ is the least upper bound of its 1-dimensional lower bounds, so Proposition 4.5(2) implies that $\tilde{c}'^* \mid c^*$, as required.

Finally, suppose that all vertices and edges in c are contained in $T_{min} \cup [*, \overline{\overline{*}}]$. Let A, B, C, D denote the four essential vertices of T_{min}, listed in the order they are numbered, from least to greatest. Let e_A, e_B, e_C, and e_D denote the edges incident with A, B, C, and D, respectively, which are in the greatest direction possible from each - namely, 2 (see the paragraphs preceding Definition 3.2). Note these are the only four edges in T_{min} which can possibly be disrespectful in a cell of $U\mathcal{D}^n T_{min}$. Since c is critical and has dimension at least 3 by assumption, c contains at least 3 edges in T_{min}, and these must be chosen from $\{e_A, e_B, e_C, e_D\}$. It follows that either e_A or e_B is in c. Let c' be the result replacing all edges in c with either endpoint, except for e_A if $e_A \in c$ or e_B if $e_A \notin c$. Let $\tilde{c}'$ be the result of moving all vertices in c' toward $*$ until they are blocked. By Lemma 4.1(4), $\tilde{c}'$ is critical.

For an arbitrary cell $\overline{c}$, a vertex $v \in \overline{c}$ is *blocked by* an edge $e \in \overline{c}$ if and only if there are no vertices of $T - \overline{c}$ which are less than v and between v and e. We claim that there are at least five vertices in $\tilde{c}'$ blocked by e_A (if $e_A \in c$) or e_B (if $e_A \notin c$). The reason is that there must be at least three edges in c, all of which are greater than or equal to e_A or e_B, respectively. As each edge must be disrespectful, each of the three edges blocks at least one vertex. When all of these strands are moved towards $*$ until they are blocked, the result is that at least five vertices are blocked in $\tilde{c}'$ by the edge e_A or e_B, respectively. This proves the claim.

It follows from this that $\tilde{c}'$ is not in the image of $\theta : U\mathcal{D}^4 T_{min} \to U\mathcal{D}^n T$, since any cell in $\theta(U\mathcal{D}^4 T_{min})$ will contain at most 4 cells from the tree T_{min}. This implies that $(\tilde{c}')^* \in \ker(\theta^*)$. By construction $[\tilde{c}'] \leq [c]$, and, since $[c]$ is the least upper bound of its 1-dimensional lower bounds, $(\tilde{c}')^* \mid c^*$ by Proposition 4.5(2). $\square$

REFERENCES

[1] Aaron Abrams. *Configuration spaces of braid groups of graphs*. PhD thesis, UC Berkeley, 2000.

[2] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[3] Kenneth S. Brown. The geometry of rewriting systems: a proof of the Anick-Groves-Squier theorem. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, volume 23 of *Math. Sci. Res. Inst. Publ.*, pages 137–163. Springer, New York, 1992.

[4] Ruth Charney and Michael W. Davis. Finite $K(\pi, 1)$s for Artin groups. In *Prospects in topology (Princeton, NJ, 1994)*, volume 138 of *Ann. of Math. Stud.*, pages 110–124. Princeton Univ. Press, Princeton, NJ, 1995.

[5] Frank Connolly and Margaret Doig. Braid groups and right angled Artin groups. *arXiv:math.GT/0411368*.

[6] John Crisp and Bert Wiest. Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups. *Algebr. Geom. Topol.*, 4:439–472, 2004.

[7] James Dugundji. *Topology*. Allyn and Bacon Inc., Boston, Mass., 1966.

[8] Michael Farber. Collision Free Motion Planning on Graphs. *arXiv:math.AT/0406361*.

[9] Michael Farber. Topological complexity of motion planning. *Discrete Comput. Geom.*, 29(2):211–221, 2003.

[10] Daniel Farley and Lucas Sabalka. Discrete Morse theory and graph braid groups. *Algebr. Geom. Topol.*, 5:1075–1109 (electronic), 2005.

[11] Daniel S. Farley. Homology of tree braid groups. *to appear in Proceedings of the Conferences on Probabilistic and Asymptotic Aspects of Group Theory and Topological Aspects of Group Theory, Edited by: R.I. Grigorchuk, M. Mihalik, M. Sapir, and Z. Sunik.*

[12] Daniel S. Farley. Homological and finiteness properties of picture groups. *Trans. Amer. Math. Soc.*, 357(9):3567–3584 (electronic), 2005.

[13] Robin Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.

[14] Robert Ghrist. Configuration spaces and braid groups on graphs in robotics. In *Knots, braids, and mapping class groups—papers dedicated to Joan S. Birman (New York, 1998)*, volume 24 of *AMS/IP Stud. Adv. Math.*, pages 29–40. Amer. Math. Soc., Providence, RI, 2001.

[15] Joseph Gubeladze. The isomorphism problem for commutative monoid rings. *J. Pure Appl. Algebra*, 129(1):35–65, 1998.

[16] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[17] Sze-tsen Hu. Isotopy invariants of topological spaces. *Proc. Roy. Soc. London. Ser. A*, 255:331–366, 1960.

[18] Carl Mautner. Private communication.

[19] M. Salvetti. Topology of the complement of real hyperplanes in $\mathbf{C}^N$. *Invent. Math.*, 88(3):603–618, 1987.

Max Planck Institute for Mathematics, D-53111 Bonn, Germany, http://www.math.uiuc.edu/~farley
*E-mail address*: `farley@math.uiuc.edu`

Department of Mathematics, University of Illinois at Urbana-Champaign, Champaign, IL 61820, USA, http://www.math.uiuc.edu/~sabalka
E-mail address: `sabalka@math.uiuc.edu`

i)

ii)

iii)